semMatch: Road Semantics-based Accurate Map Matching for Challenging Positioning Data

Heba Aly
Department of Computer Science
University of Maryland, USA
heba@cs.umd.edu

Moustafa Youssef*
Wireless Research Center
E-JUST, Egypt
moustafa.youssef@ejust.edu.eg

arXiv:1510.03533v1 [cs.CY] 13 Oct 2015

ABSTRACT

Map matching has been used to reduce the noisiness of the location estimates by aligning them to the road network on a digital map. A growing number of applications, e.g. energy-efficient localization and cellular provider side localization, depend on the availability of only sparse and coarse-grained positioning data; leading to a challenging map matching process.

In this paper, we present *semMatch*: a system that can provide accurate HMM-based map matching for challenging positioning traces. *semMatch* leverages the smartphone's inertial sensors to detect different road semantics; such as speed bumps, tunnels, and turns; and uses them in a mathematically-principled way as hints to overcome the sparse, noisy, and coarse-grained input positioning data, improving the HMM map matching accuracy and efficiency. To do that, *semMatch* applies a series of preprocessing modules to handle the noisy locations. The filtered location data is then processed by the core of *semMatch* system using a novel incremental HMM algorithm that combines a semantics-enriched digital map and the car's ambient road semantics in its estimation process.

We have evaluated *semMatch* using traces collected from different cities covering more than 150km under different harsh scenarios including coarse-grained cellular-based positioning data, sparse GPS traces with extremely low sampling rate, and noisy traces with a large number of back-and-force transitions. The results show that *semMatch* significantly outperforms traditional map matching algorithms under all scenarios, with an enhancement of at least 416% and 894% in precision and recall respectively in the most difficult cases. This highlights its promise as a next generation map matching algorithm for challenging environments.

*Moustafa Youssef is currently on sabbatical from Alexandria University, Egypt.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
SIGSPATIAL'15, November 03-06, 2015, Bellevue, WA, USA
© 2015 ACM. ISBN 978-1-4503-3967-4/15/11...$15.00
DOI: http://dx.doi.org/10.1145/2820783.2820824

Categories and Subject Descriptors

F.2.2 [**Nonnumerical Algorithms and Problems**]: Geometrical problems and computations; I.5.1 [**Computing Methodologies**]: Pattern Recognition—Models (Statistical)

Keywords

Road semantics-based map matching, Map matching for challenging positioning data, HMM-based map matching

1. INTRODUCTION

With location-based services (LBS) becoming an integral part of our everyday life, map matching; the problem of finding which road the vehicle is on given noisy input location traces; has gained attention with a large number of applications including car navigation, enriching map semantics, traffic estimation, among others. A number of map matching techniques have been proposed in literature [2, 10, 14, 18, 23, 28, 31, 37, 39], mainly designed for GPS positioning data due to its worldwide availability and relatively high accuracy. However, there are a number of situations where coarse-grained network-based location information is the only viable solution, e.g. in situations requiring energy-efficient localization, localization from the cellular provider side, and/or localization using devices that do not include a GPS chip [19–21, 40]. Therefore, a number of map matching techniques started to emerge that leverage other sensors on the phone (e.g. WiFi [34], cellular fingerprinting [35] with neighbouring cell tower information[1], and inertial sensors [15, 36]) to enhance the cellular-based localization accuracy and hence provide better map matching accuracy. Nevertheless, the sensors used are still noisy and/or consume high power, leading to coarse-grained localization; and hence lower map matching accuracy.

In this paper, we target map matching for ***challenging positioning data that is characterized by sparse, noisy, and/or coarse-grained input location estimates***. For example, cell ID-based cellular-localization systems; which estimate the user location as her associated cell tower location; are ubiquitous and energy-efficient since cellular data is available on all cell phones and using it for localization does not consume extra energy in addition to the normal phone operation as opposed to the GPS or WiFi-based systems. However, for such positioning data, the typical assumptions

[1]The majority of cell phones in the market only give access to the associated cell tower information only with no access to the neighbouring cell tower information.

(a) Cellular pos. data suffer from back-and-forth trans. The loc. estimates seq. (C1-C2-C3-C2-C4) bounces bet. C2 and C3 due to the overlapping coverage bet. the cell towers in the shaded Area A.

(b) Cellular positioning data are sparse with errors in the order of kilometers. The user driving path is fully covered by 2 cell towers (i.e. 2 location estimates).

(c) Lots of possible map matching routes and the actual route may be difficult to deduce. For example, routes 1 and 2 represent more probable map matching output by previous approaches, e.g. [25, 28].

Figure 1: Illustration of the cellular-based map matching challenges. The car moves on the red solid line. The cellular tower position is the circle and its coverage is highlighted with the dotted circle.

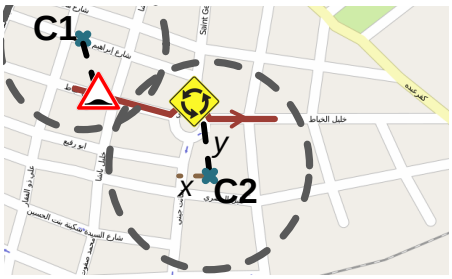

Figure 2: Traditional map matching techniques favor roads that are closer to the input location estimates. With coarse-grained accuracy, the input locations, e.g. C1 and C2, are far from the correct user trajectory. Knowledge of surrounding road semantics sensed by the energy-efficient phone sensors can reduce the ambiguity of the roads, leading to reduced complexity and high map matching accuracy.

of traditional map matching algorithms do not hold (Figure 1). Specifically, the location error is in the order of kilometers and the input location samples are highly sparse (Figure 1(b)). Moreover, the location estimates suffer from back-and-forth transitions (aka "Ping-Pong" effect [12, 22]) due to the handover between different cell-towers for call quality maintenance (even with no movement from the user) (Figure 1(a)). These stringent characteristics lead to a much harder map matching problem in terms of low quality of the input location data and the high number of candidate road segments (Figure 1(c)); affecting both the accuracy and computational complexity.

We therefore present *semMatch*, a system for accurate and efficient map matching of challenging positioning data. *semMatch* leverages the commonly available and energy-efficient inertial sensors to detect the car's ambient road semantics, such as speed bumps and tunnels, and uses them as hints to guide the map matching process and infer the car current road segment (Figure 2). At the core of *semMatch* is a *novel* incremental Hidden Markov Model (HMM) algorithm that takes the car's road surrounding information and the noise of the input data into account to enhance the accuracy of the estimated road segments and efficiently handle the increased number of candidate road segments. This HMM is combined with a number of preprocessing modules that reduce the noise in the input data.

We have implemented *semMatch* and evaluated it on driving traces collected in multiple cities. Our results show that *semMatch* can achieve a map matching F-measure of up to 97% under different challenging scenarios including coarse-grained cellular location information, highly sparse GPS traces, and traces with a lot of back-and-force transitions. Moreover, its performance is significantly better than traditional map matching techniques [25, 28] under all scenarios, reaching an enhancement of at least (416%, 894%, 232%) in (precision, recall, F-measure) respectively in the most challenging cases.

In summary, our main contributions in this paper are fourfold:

- We present the architecture of *semMatch*: a real-time map matcher for challenging positioning data. The system can accurately and efficiently map match traces with accuracy in the order of kilometers, traces with low update rates (e.g. one update every two minutes), and noisy traces with back-and-force transitions.
- We present a number of preprocessing modules that reduce the noise in the phone sensors and input locations.
- We discuss the details of a novel HMM framework that integrates a digital map with the car's ambient road semantics knowledge to accurately detect the car road segment and reduces the computational complexity.
- We implement the *semMatch* system and evaluate its performance using real driving traces covering more than 150 Km and compare it to the state-of-the-art HMM-based map matching algorithms under different harsh conditions.

The rest of the paper is organized as follows: Section 2 gives an overview of the *semMatch* system and its different components. Then, we give the details of the proposed HMM-based map matching algorithm in Section 3. Section 4 presents the experimental evaluation of the *semMatch* system. Finally, sections 5 and 6 discuss related work and conclude the paper, respectively.

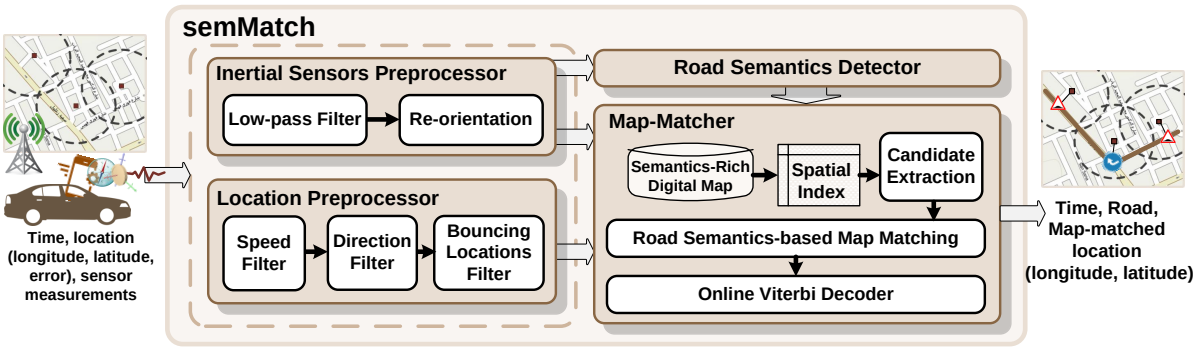

Figure 3: *semMatch* system architecture.

2. SEMMATCH OVERVIEW

In this section, we provide an overview of the *semMatch* system and the details of its different supporting modules for handling the noisiness and sparseness of the input location information. We also describe the different road semantics detected by *semMatch*. We leave the details of the core of system, *the road semantics-based map matching module*, to the next section.

2.1 Architecture Overview

Figure 3 shows *semMatch* overall architecture diagram. The input to the system is a time-stamped challenging positioning data (e.g. cellular-based localization) along with the phone's available inertial sensors (e.g. accelerometer, magnetometer and gyroscope) measurements. Each input location is represented as a longitude and latitude pair that captures the user's estimated location as well as an estimate of the localization error.

semMatch starts by preprocessing the input location points to eliminate the noise using the *Inertial Sensor Preprocessor* and the *Location Preprocessor* modules. Then, the *Road Semantics Detector* module identifies the road semantic type when the car passes by them using the inertial sensor measurements.

The preprocessed locations along with the car's current road semantics information are then passed to the *Map Matcher* module that employs a novel incremental HMM-based algorithm that integrates the hints about the car's ambient road semantics and a semantically-enriched digital map to achieve high map matching accuracy in an efficient manner. OpenStreetMaps have different tags to manually add and query the different road features such as bridges, tunnels, etc. In addition, this semantic tagging can be automated, e.g. using the Map++ system [3].

Our proposed semantic map matching algorithm contains three sub-modules: (1) Candidate Extraction and Filtering, (2) Incremental Road Semantics-based HMM Map Matching, and (3) an Online Viterbi Decoder. The *Candidate Extraction and Filtering* module determines the candidate road segments from the semantically enriched digital map. The module takes into account the error in the input location, the previous estimated user location, and the uncertainty in the semantics detection algorithm. The *Incremental Map Matching* algorithm integrates a number of modifications to the standard HMM map matching algorithm to take the detected semantic type into account as well as the road segment ambient semantics to enhance the accuracy of the estimated road segments and the map matching location on them. Finally, the *Online Viterbi Algorithm* uses dynamic programming to efficiently determine the most probable road segment. The Map Matcher outputs the matched road segment along with the car's estimated location on it (position of the last detected semantic on the road segment).

In the balance of this section, we provide the details of the Pre-processing and the Road Semantics Detector modules. We leave the details of the Map Matching module to the next section.

2.2 Preprocessing

Both the inertial sensors measurements and the input location data are known for their noisy nature [27]. These are handled by the *Inertial Sensors Preprocessor* and the *Location Preprocessor* respectively.

2.2.1 Inertial Sensors Preprocessor

This module is responsible for preprocessing the raw sensor measurements to reduce the effect of (a) phone orientation changes and (b) noise and spurious changes, e.g. sudden breaks, or small changes in the direction while moving.

First, we apply a low-pass filter to the raw sensors data using a weighted local regression filter [11]. This reduces the noise and spurious changes effect. Then, to handle the orientation changes, we transform the sensors readings from the mobile coordinate system to the world coordinate system leveraging the inertial sensors [27].

2.2.2 Location Preprocessor

For the location data, we apply a series of filters to detect false and misleading transitions and filter them out (Figure 1). In particular, we apply three filters in succession: the *Speed filter*, the *Bouncing locations filter*, and the *Direction filter*.

Speed Filter:

Typically, roads have a maximum speed limit set by traffic regulations. Also, there is a physical maximum speed that cars cannot exceed. Hence, we assume that the car will not exceed a certain speed threshold ($\nu_{\max}$). This threshold is obtained from the digital map and we also add a margin to accommodate drivers that do not strictly follow traffic rules. If a road's speed limit is missing from the map, we set the threshold to the maximum physical speed limits. If the car estimated current speed exceeds this threshold, the new location estimate is detected as an outlier and the car current location estimate remains unchanged.

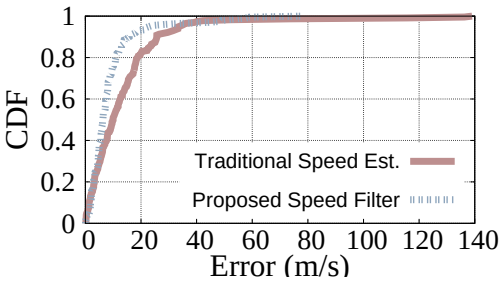

Figure 4: Using the new proposed speed filter reduces the median error by more than 36% as well as reduces the maximum error significantly.

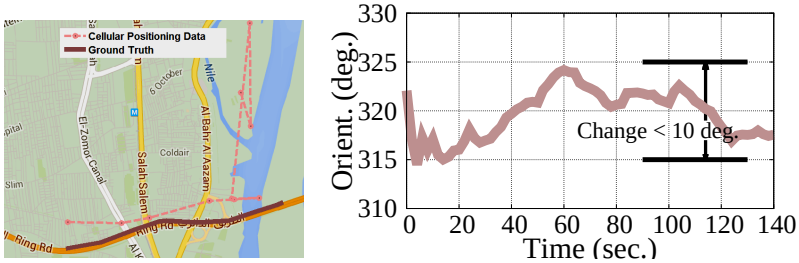

(a) Back-and-forth transitions that can indicate a false turn.

(b) The direction change is less than 10°, indicating that the user is moving in the same direction.

Figure 5: Example for back-and-forth transitions in cellular-positioning data indicating a false direction change and how *semMatch* detects this outlier using the inertial sensors readings.

To estimate the car current speed, intuitively this can be done by dividing the geodesic distance between the car's current location estimate and the new one by the difference in their time-stamps. However, due to the high error in the input locations, this speed is too noisy to use (Figure 4). Instead, we estimate the car's current speed (ν_p) by averaging the speed between the new location and a window of the preceding preprocessed locations as follows:

$$\nu_p = \frac{1}{w_s} \sum_{i=p-w_s}^{p-1} d_{i,p}/(t_i - t_p) \quad (1)$$

Where p is the index of the new location estimate, w_s is the window size, $d_{i,p}$ is the geodesic distance between locations i and p, and t_i, t_p are the time-stamps of locations i and p respectively.

Bouncing Locations Filter: To reduce the back-and-forth transitions, we apply an α-trimmed mean filter [25] on the location points. An α-trimmed filter has the advantage of handling both impulse and Gaussian noise, as compared to mean and median filters that can handle only one of them. Moreover, it is simple to implement.

The basic idea is to look at the neighbors of each point, remove 2α of the extreme neighbors, i.e. outliers, then replace the point by calculating the mean of the unfiltered neighbours. Therefore, at $\alpha = 0$, the filter works as a standard average filter while at $\alpha = 0.5$, the filter works as a median filter.

Note that to apply the filter, we need to sort the locations (loc_i). We experimented with different space-filling curves and found that applying the linear space filling curve [7] provides good accuracy while maintaining low computational time. This is intuitive as most of the time the car moves in straight lines.

Direction Filter:
Back-and-forth transitions can mislead the map matching algorithm to sense that the user made a change in her direction, i.e. as if the user took a u-turn. Therefore, to reduce this effect we apply the direction filter. It ensures that the change in the car direction is only allowed when we are sure that it is originating from an actual change in direction, not due to the input data noisiness.

To do that, we leverage the smartphone's inertial sensors to estimate the car heading direction [27]. Specifically, if a location point indicates a direction change, we confirm this change using the car estimated bearing from the inertial sensors. If the indicated direction change proved false, the car raw location remains unchanged. Figure 5 shows an example where the input cellular-location data indicates a direction change, while the inertial sensors shows that it is a false one.

Table 1: Confusion matrix for classifying different map semantics discovered from **in-vehicle** traces.

	Cat's eye	Bump	Curve	Bridge	Tunnel	Turn	U-turn	No Class
Cat's eye	**22**	0	0	0	0	0	0	5
Bump	0	**37**	0	0	0	0	0	0
Curve	0	0	**33**	0	0	0	0	0
Bridge	0	0	0	**11**	0	0	0	3
Tunnel	0	0	0	0	**15**	0	0	0
Turn	0	0	0	0	0	**55**	0	0
U-turn	0	0	0	0	0	0	**12**	0

2.3 Road Semantics Detector

The *Road Semantics Detector* module processes the smartphone's inertial sensor measurements to detect when the car passes by a road semantic and its type. In [3], we found that the smartphone sensors get affected by the different road features such as speed bumps or tunnels. For example, a car going over a speed bump will experience an up-and-down movement over a small distance which leads to a high variance in the phone's measured gravity acceleration (Figure 6(a)). Similarly, a car taking a u-turn will experience a large change in its direction (around 180°) as shown in Figure 6(b). Hence, we use these effects to detect them.

semMatch identifies a wide range of road semantics in real-time using a decision-tree classifier (Figure 7) including **tunnels, curves, bridges, turns, speed bumps, u-turns, and cat's eyes**. We refer the readers to [3] for more information about the features we used to identify the different road semantics using the smartphone's sensors. Table 1 shows the confusion matrix between the different classes. The different road semantics can be identified with high accuracy due to their unique signatures. However, to compensate for cases not observed during the classifier training, we

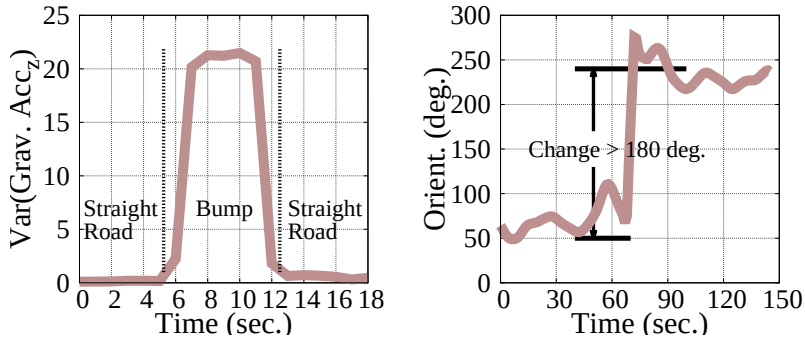

(a) Going over a speed bump leads to a large change in the phone's gravity acceleration.

(b) Taking a u-turn leads to a change in the user direction (i.e. phone orientation) around 180°.

Figure 6: Passing by the different road semantics (e.g. speed bumps and u-turns) has unique effects on the smartphone's sensors measurements.

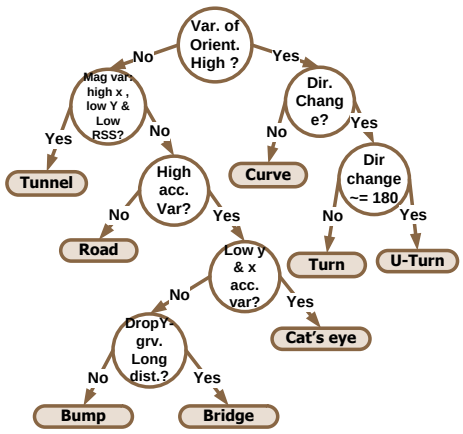

Figure 7: Decision tree used by *semMatch* to identify the different road semantics in real-time.

add a small error to every cell in the matrix to account for further confusion.

3. SEMANTIC MAP MATCHING

Through this section, we provide the details of our novel road semantics-based *Map Matcher* module. We first start by providing an overview on the new model and its differences from previous HMM algorithms. Then, we provide the mathematical model and the details of the HMM model components.

3.1 Overview

In traditional map matching using GPS traces, e.g. [28], the input to the system is a trace of the noisy GPS locations and the output is the corresponding sequence of road segments the car passed by. Extending this directly to the coarse-grained, noisy, and sparse cellular location information (Figure 1) leads to poor map matching accuracy with high complexity as explained in Section 1. *semMatch* solves these challenges by leveraging the road semantics the car passes by to reduce ambiguity and hence enhances both the map matching accuracy and running time. Specifically, the input to the map matcher module is a sequence of ordered triples in the form of **(Coarse-grained location, Estimated location error, Semantic type)** representing the road semantics detected by the phone sensors during the car movement at a specific coarse-grained input location. Note that due to the rich number of semantics detected by *semMatch* (as described in Section 2.3), the input data to *semMatch* is much more frequent than, e.g. a sequence of unique cell IDs only, and hence should lead to better accuracy and fewer road segments candidates as described in details in the rest of this section.

3.2 Mathematical Model

A HMM can be represented as $\lambda = (S, V, A, B, \pi)$ [32] where:

- $S = \{s_1, ..., s_N\}$ is the set of possible states and $N = |S|$. In our case, each state represents a road semantic landmark (e.g. a speed bump) on a specific road segment as shown in Figure 8. Therefore, a state s_i is represented by the ordered triples **(Road segment ID (ID), Road semantic landmark type (type), Road semantic landmark location (loc))**, where the "Road semantic landmark location" represents the coordinates of the semantic landmark on the given road segment. ***Note that using this representation, each physical road segment is split into a number of states by the different semantic landmarks on this segment*** (Figure 8).

- $V = \{v_1, ..., v_M\}$ is the set of observations from the model and $M = |V|$. In our case, each observation is a detected road semantic type using the phone inertial sensors and is represented by an ordered quadruple in the form of **(Coarse-grained location (loc), Estimated location error (err), Estimated heading direction (θ), Detected semantic type (type))**.

- $A = \{a_j(k)\}$ is the observation symbol probability distribution in state j, where $a_j(k) = P[z_t = v_k | q_t = s_j], i < j < N, 1 < k < M$ and z_t and q_t are the observation and state at time t, respectively. In other words, this represents the probability of observing a certain road semantic type given the car is at a specific state, i.e. road segment with a certain road semantic landmark.

- $B = \{b_{ij}\}$ is the state transition probability distribution, where $b_{ij} = P[q_{t+1} = s_j | q_t = s_i], i, j < N$.

- $\pi = \{\pi_i\}$ is the initial state distribution, where $\pi_i = P[q_1 = s_i]$.

Therefore, the problem becomes, given a sequence of observations $Z = (z_1, ..., z_T)$, where T is a system parameter, and each $z_i \in V, 1 < i < T$, we want to find the most probable sequence of semantic road segments (states) $Q = (q_1^*, ..., q_T^*)$, where each $q_i^* \in S, 1 < i < T$. In the rest of this section, we give the details of how *semMatch* models these different components.

3.3 Observation Probability Distribution (A)

Since there can be ambiguity in the detected semantics due to the noisy phone sensors, e.g. a cat's eye detected as a speed bump, the probability of the phone detecting a given road semantic Y (e.g. a speed bump) given a specific road

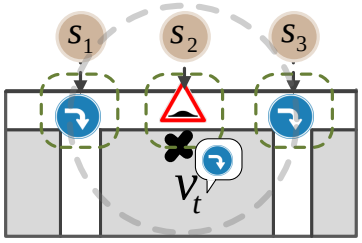

Figure 8: Proposed road semantics-based HMM model. A state (s_i) represents a road segment semantic landmark and is modelled as an ordered triples <Physical road segment ID, Landmark type (e.g. bump in state s_2), Landmark location on the road>. The user observations (v_t's) are defined by the estimated location (x mark), the detected semantic type (e.g. turn in v_t), and the estimated error (dotted circle). The observation probability ($p(v_t|s_i)$) is a function of both the road landmark's location and type and the observation's location and type.

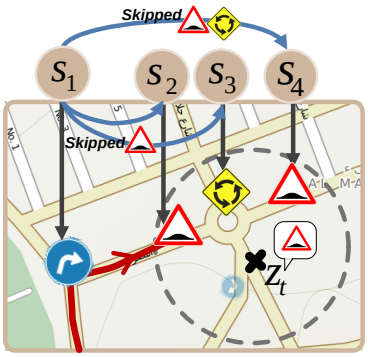

Figure 9: Transition probabilities calculation from state s_1 to states s_2, s_3, and s_4. The transition probability is a function of 1) the difference in the heading between the hidden states' road segments on the map and the car change of heading as sensed by the phone sensors and 2) the skipped landmarks on the map. States s_3 and s_4 are penalized due to skipping one or more road semantics.

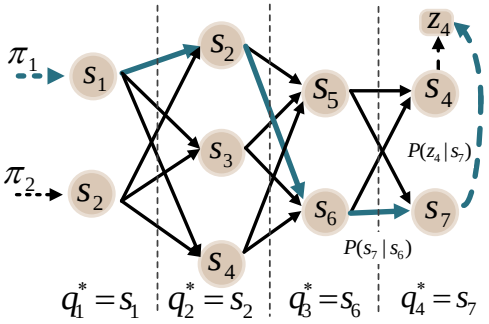

Figure 10: Illustration of the proposed HMM model Viterbi decoding: States are the graph vertices, states transition and observation probabilities are drawn as links in the graph. The initial states probabilities (π_i's) are estimated for the first set of states in the current sequence. The Viterbi decoded sequence (Q) is colored in blue. The user observation is drawn here only at time $t = 4$ due to the incremental nature of the algorithm.

segment semantic landmark X (e.g. a cat's eye) should be a function of two factors (Figure 8): 1) the distance between the location of the car where the semantic Y was detected and the location of the semantic landmark X on the road segment. The higher this distance, the lower the probability that should be assigned to a given road segment. 2) the probability of the phone sensors detecting type Y when the car is passing by a semantic of type X. This is captured by the false positive and false negative rates of detecting the different semantics as in Section 2.3.

semMatch models the *first factor* as a Gaussian distribution:

$$f_1 = \frac{1}{\sqrt{2\pi}z.\sigma} e^{-0.5\left(\frac{\text{dist}(z.loc, s_i.loc)}{z.\sigma}\right)^2} \quad (2)$$

Where dist($z.loc, s_i.loc$) is the geodesic distance between the observation z and the road segment semantic, i.e. hidden state, s_i. $z.\sigma$ is the standard deviation of a Gaussian random variable that corresponds to the error in the observed input location.

For the *second factor*, we model it using the *Road Semantics Detector* confusion matrix to estimate the probability of observing a semantic type $z.type = Y$ at time t conditioned on the hidden state s_i road semantic landmark type ($s_i.type = X$):

$$f_2 = p(z.type|s_i.type) \quad (3)$$

Hence, the final observation probability, $p(z|s_i)$, is modeled as:

$$p(z|s_i) = p(z.type|s_i.type) \frac{1}{\sqrt{2\pi}z.\sigma} e^{-0.5\left(\frac{\text{dist}(z.loc, s_i.loc)}{z.\sigma}\right)^2} \quad (4)$$

3.4 Transition Probability Distribution (B)

The transition probability is the probability of moving to the next state $s_{j,t}$ given the current state is $s_{i,t-1}$. Intuitively, for a probable transition between two road segment landmarks, the change in the car orientation (as detected by the phone sensors) between the two segments should match the change of the orientation in the digital map for the same two segments.

Therefore, *semMatch* models this intuition by a function of the heading change in the input observations ($\Delta\theta_z$) and the heading change between the hidden states pair ($\Delta\theta_s$) using a Gaussian distribution as:

$$f(\Delta\theta_z, \Delta\theta_s) = \frac{1}{\sqrt{2\pi}\sigma_h} e^{-0.5\left(\frac{|\Delta\theta_z - \Delta\theta_s|}{\sigma_h}\right)^2} \quad (5)$$

Where σ_h is the estimated standard deviation of the heading change error. The more similar the change in heading direction, the higher the transition probability. However, in some cases, due to the high error in the input location, we can have more than one hidden state with a similar semantic type (Figure 9). However, misleading semantics will lead to skipping road segment landmarks. Therefore, to capture this in the transition probability, we penalize transitions that skip one or more road semantics. The final transition probability in *semMatch* is estimated as:

$$p(s_j|s_i) = \frac{1}{\sqrt{2\pi}\sigma_h} e^{-0.5\left(\frac{|\Delta\theta_z - \Delta\theta_s|}{\sigma_h}\right)^2} \prod_{l \in L} p(\bar{l}|l) \quad (6)$$

Where L is the set of skipped road semantic segments between the states pair s_j and s_i. For a semantic $l \in L$, the probability $p(\bar{l}|l)$ represents the probability of passing by the semantic l and not detecting it. We model this probability using the false negative value of the semantic l computed from the confusion matrix of our *Road Semantics Detector* module (Table 1).

To estimate σ_h, we use the ground-truth data. Note that, unlike the input positioning data which has highly dynamic and variable accuracy, the sensors heading accuracy has a more consistent accuracy outdoors [4]. Hence, we use a fixed σ_h calculated based on the Median Absolute Deviation (MAD) of the direction changes in our ground truth data, which is a robust estimator of σ_h [4, 28]:

$$\sigma_h = 1.4826 \times median(|\Delta\theta_z - \Delta\theta_s|) \quad (7)$$

3.5 Initial State Distribution

At each time instance, there is a large number of candidate road semantic segments that the map matcher can work with, which affects both the system accuracy and running time negatively. To reduce the search space, the *Candidate Extraction and Filtering* module leverages the the current coarse-grained cellular input location estimate (*loc*), the associated estimated error in this location (*err*), and the last estimated candidate road segments. In particular, the module works in two steps: First, it extracts the candidate road semantics from the digital map that fall inside the circle centered at *loc* with radius *err*. To speed up this process, it builds an R-tree spatial index [17] on all possible road semantic segments in the road network. Second, it remove candidate road semantic segments that are not connected to any of the candidate road segments from the previous estimation step.

Initially, all road semantics landmarks selected by the *Candidate Extraction and Filtering* module are assigned weights based on their semantic type as:

$$\pi_{i,1} = p(z_1.type|s_i.type) \tag{8}$$

which is obtained from the semantic types confusion matrix. This is intuitive as if the user observed a bump, the states with a bump as their landmark should have initial probability higher than states with other semantic types (e.g. tunnels).

The initial state distribution is re-estimated after each step in our semantics HMM for the new states as the product of the old initial state distribution and the state's transition probability as:

$$\pi_{i,t} = \sum_{j \in S_{t-1}} \pi_{j,t-1} p(s_{i,t}|s_{j,t-1}) \tag{9}$$

3.6 Optimal State Sequence Estimation

The Viterbi decoder algorithm aims to determine the most probable hidden states sequence based on the estimated HMM probabilities. The selected hidden states represent the car's traversed road semantic segments.

For *semMatch* to operate in real-time, it cannot wait till the whole sequence is available. Hence, in an incremental manner, *semMatch* uses a sliding window on the *Map Matcher* module input locations. Every time a new location is introduced, the HMM parameters are calculated for the new introduced location and the associated candidate states. The online Viterbi algorithm [9, 33] is then applied to compute the maximum likelihood sequence of hidden states for the current window using dynamic programming by extending the current solution with the estimated HMM parameters to get the best path (Figure 10). Also, the initial states probability is updated as the sliding window moves.

Algorithm 1 summarizes the *semMatch* map matching algorithm.

Algorithm 1 *semMatch* Map Matching Algorithm

Input: z_t ▷ User observed point at time t
Output: loc, id ▷ map matched user location and road segment id
Let $S = \{s_1, s_2, ..., s_N\}$ be the hidden states set.
Let $S_{t-1}[]$ be the previous hidden states vector set;
Let $\Pi_{t-1}[]$ be vector set for the states' initial probability at time $t-1$
Procedure:
$S_t \leftarrow$ ***ExtractCandidates***(S, z_t)
▷ Extract candidates road semantics
for each $s_i \in S_t$ **do**
ob[i] $\leftarrow p(z_t.type|s_i.type)\frac{1}{\sqrt{2\pi}z_t.\sigma}e^{-0.5\left(\frac{\text{dist}(z_t.loc, s_i.loc)}{z_t.\sigma}\right)^2}$
▷ Calculate Observation Probability
if ***isEmpty***(S_{t-1}) **then**
$\Pi_{t-1}[i] \leftarrow p(z_t.type|s_i.type)$
else
for each $s_j \in S_{t-1}$ **do**
$L \leftarrow$ ***SemanticsSearch***(S, s_j, s_i)
▷ Search for semantics between s_i and s_j
FN $\leftarrow 1$
for each $l \in L$ **do**
FN $\leftarrow$ FN * $p(\bar{l}|l)$
end for
tr[i][j] $\leftarrow \frac{1}{\sqrt{2\pi}\sigma_h}e^{-0.5\left(\frac{|\Delta\theta_z - \Delta\theta_s|}{\sigma_h}\right)^2} \times$ FN
▷ Calculate Transition Probability
end for
end if
end for
Q,$\Pi_{t-1} \leftarrow$ ***OViterbiDecoder***$(\Pi_{t-1}$***[],tr[][],ob[])*** ▷ Apply Online Videtrbi Decoder to get the most probable sequence and the updated initial probability
output $q_t.loc, q_t.ID$
▷ Location and road id of last state in the sequence Q

4. EXPERIMENTAL EVALUATION

In this section, we show our evaluation for the performance of *semMatch*. We stress-tested the system using real challenging positioning data traces including: (1) Cellular-based traces with accuracy in the order of kilometers, (2) Network-based (Cellular+WiFi) with a large number of back-and-force transitions, and (3) Low-sampling rate GPS-based trajectories (one sample every two minutes). Our data consists of over 150 km of car driving traces collected in two different cities. The data was collected using different Android phones including Samsung Galaxy S3 Mini, S5, and LG Nexus 4. We also used OpenStreetMaps for the road network information.

To get the cellular-based location estimates, we use the Google Android Location API **with WiFi turned off**. This way, the returned positioning estimates are based on the associated cell tower information only since the phones we used does not provide the neighboring cell towers information. The cellular-positioning data had *an average accuracy of 1.9 km and density of 1.4 location estimate per km.* To get the network-based location estimates, we use the Google Android Location API **with WiFi turned on**. The network-positioning data had *an average accuracy of 162m and density of 6.8 location estimate per km.* For the low-sampling rate GPS-data, we used the cellphones' GPS sensor with a sampling rate of one sample every two minutes. The GPS-positioning data had *an average accuracy of 19m and density of 0.6 location estimate per km.* Cellular and network-based traces are noisier than GPS-based traces and has much higher errors. However, for energy efficiency, the GPS is sampled at low sampling rates; leading to extremely sparse positioning data.

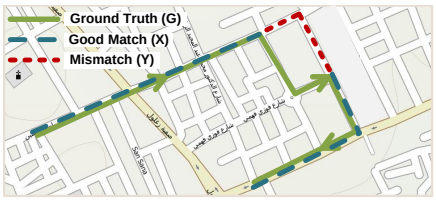

Figure 11: An example showing the definition of precision $(X/(X+Y))$ and recall (X/G).

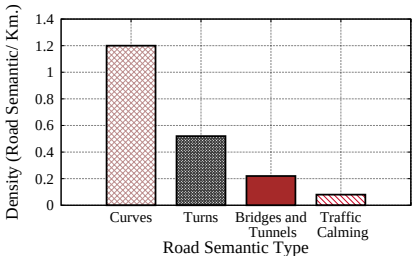

Figure 12: Density of the different road semantics in our trajectory traces.

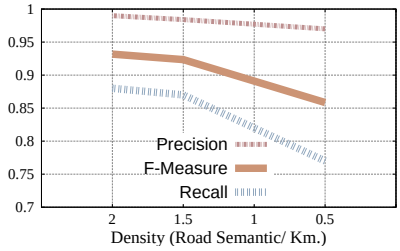

Figure 13: Performance of *semMatch* under different road semantics density.

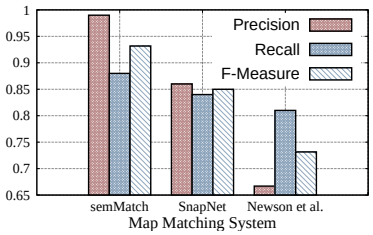

(a) Cellular Pos. Data (average pos. accuracy=1.9km, avg. update rate= 1.4 loc. estimate per km)

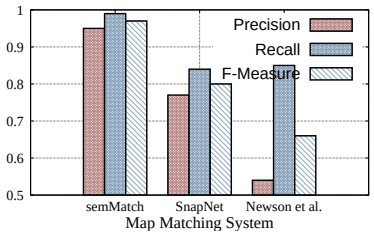

(b) Network Pos. Data "Cellular+WiFi" (a large number of back-and-forth transitions)

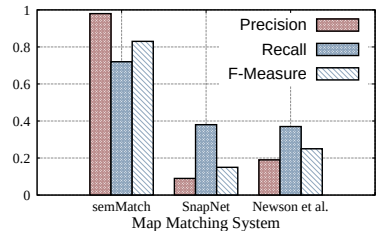

(c) Low Sampling Rate GPS Positioning Data (1 sample every 2 mins)

Figure 14: Performance of *semMatch* as compared to SnapNet [25] and Newson et al [28] under different **challenging positioning data** classes.

For the ground truth, we collected GPS/GlONASS-based location data at high sampling rate using an external receiver. Then, we map matched the traces using offline HMM-based map matching. It has been shown in literature that GPS traces with high sampling rate (>1 sample per second) can be map matched to reconstruct the driving path [25,28].

The rest of this section is organized as follows: We start by describing our evaluation metrics. After that, we show the effect of the road semantics density on *semMatch*'s performance. Then, we quantify the *semMatch* map matching performance as compared to two map matching techniques from literature. Finally, we show the power consumption overhead when using *semMatch*.

4.1 Evaluation Metrics

To evaluate *semMatch* map matching accuracy, we find the common matched road segment sequence between the output map matched trajectory and the ground-truth trajectory. Then, we compute the map matching *precision and recall* using this common sequence (Figure 11). The map matching precision shows how accurate the map matched trace is; it is defined as the ratio between the distance traversed on the matching sequence X and the total distance of the map matched trajectory $(X+Y)$. On the other hand, the map matching recall shows the actual trajectory recovered ratio; it is defined as the ratio between the distance traversed on the matching sequence X and the total distance of the ground truth trajectory G.

More formally, we can define the metrics as:

$$\begin{aligned} Precision &= \frac{\text{Total distance of common matching sequence}}{\text{Total distance of output trace}} \\ &= \frac{X}{X+Y} \end{aligned} \tag{10}$$

$$\begin{aligned} Recall &= \frac{\text{Total distance of common matching sequence}}{\text{Total distance of ground truth}} \\ &= \frac{X}{G} \end{aligned} \tag{11}$$

We also calculate the map matching F-Measure (F_1 score) which is the harmonic mean of the precision and recall:

$$F\text{-}Measure = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{12}$$

4.2 Effect of Road Semantics Density on Accuracy

Figure 12 shows the density of the different road semantics encountered in our trajectories. On average, the density of the different road semantics is 2 semantics per km; this is higher than the cellular location update rate by 43%, which allows *semMatch* to achieve better map matching accuracy.

Figure 13 shows the effect of the road semantics density on the map matching precision and recall. The figure shows the higher the density, the higher the system accuracy. *semMatch* can achieve an F-measure of 93% using the typical road semantics density.

4.3 Comparison with Other Systems

In this section, we compare the performance of *semMatch* in terms of precision, recall, and F-measure to two traditional HMM map matching algorithms: the one by Newson et al [28] and the SnapNet system [25]. Newson et al [28] use the standard HMM-based algorithm to map match noisy and relatively sparse (up to one sample every 30 sec.) GPS traces. On the other hand, SnapNet targets map matching noisy cellular-positioning data. They apply several preprocessing stages based on the road network topology and assumed that the users favor major roads to overcome the imposed cellular-based challenges.

We used different scenarios of challenging positioning data:

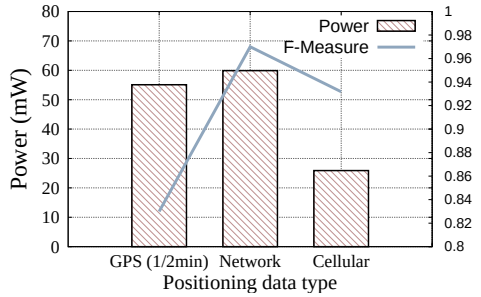

Figure 15: *semMatch* power consumption for the different positioning data types (low sampling rate GPS, network, and cellular) and the corresponding map matching F-measure.

(1) Cellular-based data (avg. position accuracy=1.9 km, avg. location update= 1.4 location estimate per km), (2) Network-based data (WiFi+Cellular) (noisy data with a lot of back-and-forth transitions), and (3) GPS with low sampling rate (1 sample every 2 minutes).

Figure 14 shows that *semMatch* has the ***significantly better accuracy compared to the other two techniques under all scenarios***. For the most challenging case of extremely low sampling rate of one location estimate every two minutes, *semMatch*] leads to enhancement in the (precision, recall, F-measure) of (989%, 894%, 232%) over the Newson et al method [28] and (416%, 946%, 453%) over the SnapNet system. This highlights the advantage of incorporating the road semantics in the map matching process under different harsh conditions.

4.4 Power Consumption

semMatch collects inertial sensor measurements along with the positioning data to detect the different road semantics and improve the map matching accuracy. Figure 15 shows the system's power consumption along with the map matching F-measure. The figure shows that *semMatch* could achieve high map matching accuracy for the different challenging positioning data types with a low-energy footprint. Note that *semMatch* depends on the smartphone's inertial sensors to detect the road semantics. Inertial sensors have low power consumption and are already running for other purposes (e.g. screen orientation detection). Hence, using them for road semantics detection consumes zero extra energy.

5. RELATED WORK

5.1 Map Matching Techniques

Due to its importance in many applications, map matching algorithms have gained due attention from researchers [2, 10, 14, 18, 23, 25, 28, 31, 39]. These algorithms employ different techniques including geometric, topological, and probabilistic ones [30]. Geometry-based map matching algorithms utilize the geometry and shape of the input trace and the road network. For example, in [38], authors map match the input location point to its closest arc (sequence of road map nodes). Similarly, in [29, 38], authors use geometric curve to curve matching; where they map match a series of points simultaneously to the closest arc. These techniques consider only the shape of the links regardless of their connectivity, leading to false transitions. Moreover, while they could provide good map matching accuracy for clean and dense GPS trajectories, they are unsuitable for noisy and sparse positioning data. On the other hand, topology-based map matching algorithms utilize the road segments topological information, e.g. connectivity and contiguity, along with the input trace and the road network shapes [8, 14, 31]. For example, in [31], authors use various similarity criteria (e.g. vehicle speed, closeness) between the road network geometry and derived GPS navigation data to map match the user trajectory to the road network. While topology-based algorithms improve upon geometry-based ones by considering road topology and characteristics, they are still vulnerable to the highly erroneous and sparse data.

Finally, probabilistic map matching algorithms employ probabilistic frameworks such as HMM or particle filters to map match the input user trajectories [16, 18, 28]. For example, in [18, 28], HMM-based techniques give relatively good accuracy for the relatively noisy GPS trajectories with up to 30 seconds delay. However, highly noisy and sparse data are much harder to map match as shown in Figure 1. Hence, their accuracy degrades significantly, up to 0.25 map matching F-Measure for low sampling rate GPS trajectories (as we quantified in Section 4).

Recently, few systems were proposed to solve the map matching problem for coarse-grained network-based location data [15, 25, 34–36]. Usually, these techniques try to enhance the accuracy of the input location data by leveraging other sensors. For example, VTrack [34] builds an HMM-based map matching scheme leveraging the WiFi data to handle the inaccuracy of cellular location information. CTrack [35] alternatively uses *war-driving training* data in addition to inertial sensors to reduce the inaccuracy of cellular-based locations. Collecting war-driving data is an expensive and time-consuming process. Moreover, the cellular war-driven data typically has heterogeneity problems and need to be updated from-time-to-time. The AutoWitness [15] and WheelLoc [36] systems leverage inertial sensors and dead-reckoning to reduce the inaccuracy of cellular-based localization. However, using the accelerometer and compass for dead-reckoning leads to fast error accumulation for their estimated trajectories [1, 5, 6, 27]. SnapNet [25] targets cell-Id based positioning data by leveraging information from the digital map including road types and speed; they favor major roads and staying on the same road. Hence, the system performance would degrade when the user traverses side roads which are typically a part of any trip.

semMatch is ***unique in employing a general road semantics-based map matching framework*** that leverages a wide range of road semantics landmarks such as tunnels, curves and bridges to improve the map matching accuracy through a novel HMM-based model. *semMatch* takes advantage of the inertial sensors ***energy efficiency*** to detect the various road semantics and avoid their dead-reckoning issues. This leads to ***enhancements in the map matching F-Measure of at least*** 232% when compared to traditional and state-of-the-art HMM map matching algorithms [25, 28] as quantified in Section 4.

5.2 Road Semantics Detection

Inertial sensors have been used in literature for monitoring road conditions, e.g. [13, 24, 26]. In [13, 26], external accelerometers were used to detect potholes and traffic conditions while in [24], the smartphone's accelerometer was

used to detect potholes. All these systems use GPS to localize the road problems. In [3], we extended these systems by leveraging different smartphone's sensors to detect a wide range of outdoor road semantics with the goal of automatically enriching the current digital maps. *semMatch* builds on these road-semantics enriched digital-map to provide accurate map matching for challenging positioning data.

6. CONCLUSIONS

We presented *semMatch*: a novel road semantics-based map matcher for the challenging cellular-based trajectories. *semMatch* leverages the smartphone's inertial sensors to detect different road semantics and uses them in a mathematically-principled way to improve the accuracy and efficiency of the HMM map matching algorithm. We provided the *semMatch*'s system architecture and our different preprocessing modules that help it reduce the noise in the input data. Moreover, we presented our novel incremental road semantics-based HMM algorithm in detail.

Evaluation of *semMatch* on different classes of challenging positioning traces collected from two different cities covering more than 150km shows that *semMatch* can significantly outperform the traditional map matching algorithms in all scenarios reaching an F-Measure up to 97%. This maps to at least 232% enhancement in the F-measure in the most challenging case. In addition, *semMatch* has a low energy footprint on the scarce phone battery.

Acknowledgment

This work is supported in part by a grant from the Egyptian Information Technology Industry Development Agency (ITIDA). We also sincerely thank *Reham Mohamed* for her valuable contributions and suggestions to the *semMatch* algorithm.

7. REFERENCES

[1] H. Abdelnasir, R. Samir, A. Elgohary, M. Elzantot, H. Wang, S. Sen, R. Choudhury, and M. Youssef. SemanticSLAM: Using environment landmarks for unsupervised indoor localization. *Mobile Computing, IEEE Transactions on*, Accepted for publication.

[2] H. Alt, A. Efrat, G. Rote, and C. Wenk. Matching planar maps. In *Proc. of the ACM-SIAM symp. on Discrete algorithms*. Society for Industrial and Applied Mathematics, 2003.

[3] H. Aly, A. Basalamah, and M. Youssef. Map++: A crowd-sensing system for automatic map semantics identification. In *SECON*. IEEE, 2014.

[4] H. Aly, A. Basalamah, and M. Youssef. Lanequest: An accurate and energy-efficient lane detection system. *IEEE PerCom*, 2015.

[5] H. Aly and M. Youssef. Dejavu: an accurate energy-efficient outdoor localization system. In *SIGSPATIAL*. ACM, 2013.

[6] M. Alzantot and M. Youssef. Uptime: Ubiquitous pedestrian tracking using mobile phones. In *WCNC*, pages 3204–3209, 2012.

[7] M. Bader. *Space-Filling Curves - An Introduction with Applications in Scientific Computing*, volume 9 of *Texts in Comp. Sc. and Eng.* Springer-Verlag, 2013.

[8] C. A. Blazquez and A. P. Vonderohe. Simple map-matching algorithm applied to int. winter maintenance vehicle data. *Transp. Research Record: Journal of the Transp. Research Board*, 2005.

[9] J. Bloit and X. Rodet. Short-time Viterbi for online HMM decoding: Evaluation on a real-time phone recognition task. In *ICASSP*. IEEE, 2008.

[10] S. Brakatsoulas, D. Pfoser, R. Salas, and C. Wenk. On map-matching vehicle tracking data. In *VLDB 2005*.

[11] W. S. Cleveland and S. J. Devlin. Locally weighted regression: An approach to regression analysis by local fitting. *Journal of the American Statistical Assoc.*, 1988.

[12] N. Ekiz and T. Salih at al. An overview of handoff techniques in cellular networks. *Inter. journal of information technology*, 2005.

[13] J. Eriksson and L. Girod et al. The pothole patrol: Using a mobile sensor network for road surface monitoring. In *MobiSys*, pages 29–39. ACM, 2008.

[14] J. S. Greenfeld. Matching GPS observations to locations on a digital map. In *National Research Council (US). Transp. Research Board. Meeting*, 2002.

[15] S. Guha and K. Plarre at al. Autowitness: locating and tracking stolen property while tolerating GPS and radio outages. *TOSN*, 2012.

[16] F. Gustafsson and F. Gunnarsson et al. Particle filters for positioning, nav., and tracking. *IEEE Trans. on Signal Processing*, 2002.

[17] A. Guttman. *R-trees: a dynamic index structure for spatial searching*. ACM, 1984.

[18] B. Hummel. Map matching for vehicle guidance. *Dynamic and Mobile GIS: Investigating Space and Time*, 2006.

[19] M. Ibrahim and M. Youssef. Cellsense: a probabilistic RSSI-based GSM positioning system. In *GLOBECOM*, pages 1–5. IEEE, 2010.

[20] M. Ibrahim and M. Youssef. A hidden markov model for localization using low-end GSM cell phones. In *ICC*. IEEE, 2011.

[21] M. Ibrahim and M. Youssef. Cellsense: An accurate energy-efficient gsm positioning system. *IEEE TVT*, 61:286–296, 2012.

[22] S. Lee and K. Son et al. Base station association in wireless cellular networks: An emulation based approach. *IEEE Trans. on Wireless Communications*, 2012.

[23] X. Li, M. Li, W. Shu, and M. Wu. A practical map-matching algorithm for GPS-based vehicular networks in shanghai urban area. In *CCWMSN07*. IET, 2007.

[24] A. Mednis, G. Strazdins, R. Zviedris, G. Kanonirs, and L. Selavo. Real time pothole detection using android smartphones with accelerometers. In *DCOSS*. IEEE, 2011.

[25] R. Mohamed, H. Aly, and M. Youssef. Accurate and efficient map matching for challenging environments. In *SIGSPATIAL*. ACM, 2014.

[26] P. Mohan and V. N. et al. Padmanabhan. Nericell: rich monitoring of road and traffic conditions using mobile smartphones. In *SenSys*. ACM, 2008.

[27] N. Mohssen, R. Momtaz, and H. Aly et al. It's the

human that matters: accurate user orientation estimation for mobile computing applications. In *MobiQuitous 2014*.

[28] P. Newson and J. Krumm. Hidden markov map matching through noise and sparseness. In *SIGSPATIAL*. ACM, 2009.

[29] B. P. Phuyal. Method and use of aggregated dead reckoning sensor and GPS data for map matching. In *ION GPS*, 2001.

[30] M. A. Quddus, W. Y. Ochieng, and R. B. Noland. Current map-matching algorithms for transport applications: State-of-the art and future research directions. *Transp. Research Part C: Emerging Tech.*, 2007.

[31] M. A. Quddus, W. Y. Ochieng, L. Zhao, and R. B. Noland. A general map matching algorithm for transport telematics applications. *GPS solutions*, 2003.

[32] L. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2), Feb 1989.

[33] R. Šrámek, B. Brejová, and T. Vinař. On-line viterbi algorithm and its relationship to random walks. *arXiv preprint arXiv:0704.0062*, 2007.

[34] A. Thiagarajan et al. VTrack: accurate, energy-aware road traffic delay estimation using mobile phones. In *SenSys*. ACM, 2009.

[35] A. Thiagarajan et al. Accurate, low-energy trajectory mapping for mobile devices. In *Proc. of conf. on Networked Sys. design and implementation*. USENIX, 2011.

[36] H. Wang and Z. Wang et al. Wheelloc: Enabling continuous loc. service on mobile phone for outdoor scenarios. In *INFOCOM*. IEEE, 2013.

[37] G. Wang et al. Eddy: an error-bounded delay-bounded real-time map matching algorithm using HMM and online Viterbi decoder. In *SIGSPATIAL*. ACM, 2014.

[38] C. E. White, D. Bernstein, and A. L. Kornhauser. Some map matching algorithms for personal nav. assistants. *Transp. Research Part C: Emerging Techn.*, 2000.

[39] J.-s. Yang et al. The map matching algorithm of GPS data with relatively long polling time intervals. *J. of the Eastern Asia Society for Transp. Studies*, 2005.

[40] M. Youssef, M. A. Yosef, and M. El-Derini. GAC: energy-efficient hybrid GPS-accelerometer-compass GSM localization. In *GLOBECOM*. IEEE, 2010.